

FORMALIZING REASONS, OUGHTS, AND REQUIREMENTS

ROBERT MULLINS

The University of Queensland

Reasons-based accounts of our normative conclusions face difficulties in distinguishing between what ought to be done and what is required. This article addresses this problem from a formal perspective. I introduce a rudimentary formalization of a reasons-based account and demonstrate that that the model faces difficulties in accounting for the distinction between oughts and requirements. I briefly critique attempts to distinguish between oughts and requirements by appealing to a difference in strength or weight of reasons. I then present a formalized reasons-based account of permissions, oughts and requirements. The model exploits Joshua Gert (2004; 2007) and Patricia Greenspan's (2005; 2007; 2010) suggestion that some reasons perform a purely justificatory function. I show that the model preserves the standard entailment relationships between requirements, oughts and permissions.

1. Introduction

Many philosophers accept that what we ought to do is explained by our practical reasons. Variants of this reasons-based view are held by T.M. Scanlon (1998), Joseph Raz (1999a), Jonathan Dancy (2004), Mark Schroeder (2007), and Derek Parfit (2011). The development of formal models has made it possible to explore the relationship between reasons and conclusions about what we ought to do with greater precision (Horty 2012; Nair 2016; for a survey see Nair & Horty 2018). The purpose of this article is to explore the merits of such a model if it is extended to a broader set of normative conclusions. In particular, I consider whether the model can be extended to distinguish between what is permissible, what we ought to do, and what is required.

Contact: Robert Mullins <robertlyonmullins@gmail.com>

Reasons-based accounts of our normative concepts face difficulties in distinguishing between what we ought to do and what is required or impermissible (see especially Bedke 2011 and Snedegar 2016). A related challenge to reasons-based accounts emerges from what Raz refers to as the ‘basic belief’ that in our practical lives it is often the case that many different options are rationally intelligible, so that it would be permissible to pursue them but it ‘would not be against reason to avoid any one of them’ (1999b: 100). In some contexts, our reasons support the conclusion that certain actions are rationally intelligible but not required. Put in another way, there are some actions that we ought to perform but that it would not be impermissible for us not to perform. This sort of optionality threatens more straightforward attempts to offer reasons-based accounts of our normative conclusions. Perhaps unsurprisingly, formal approaches to modelling the relation between reasons, oughts and requirements face similar difficulties.

In order to address these problems from a formal perspective, in Section 2 I introduce a basic model of the relationship between reasons and normative conclusions. In Section 3 I explore its application to the idea that reasons explain statements of requirement and permission, as well as ‘oughts’.¹ In Section 4, I briefly critique the idea that a difference in strength or weight of reasons can explain the difference between oughts and requirements. In Section 5, I conclude by proposing a revised model that addresses these problems, based on Joshua Gert (2004; 2007) and Patricia Greenspan’s (2005; 2007; 2010) proposal that some reasons fulfil a purely justificatory function. The model provides a formal account of the relationship between reasons, permissions, oughts, and requirements. What we are required to do (and therefore what it would be impermissible for us not to do) is explained by our best requiring reasons. What we ought to do is explained by our best justifying reasons.

2. The Basic Model: Reasons, Oughts, and Requirements

The basic model relied upon here to explore the relationship between reasons and requirements is rudimentary version of an ‘imperatival’ or default-based model.² Variants of this model have been used to great effect by John Horty

1. I will follow what is now the standard philosophical practice of using the nominalization ‘ought’ (and its plural ‘oughts’) to refer to conclusions about what ought to occur.

2. The approach has its origins in van Fraassen (1973). Jörg Hansen (2004; 2005) develops his models as an interpretation of the relationship between commands or imperatives and deontic conclusions. Horty (2003; 2012) develops a deontic interpretation of Reiter’s default logic in terms of reasons, relating his approach back to van Fraassen’s initial proposal. For a logical overview of these approaches see Goble (2013: §§ 4.4, 6.3).

(2012) and Shyam Nair (2016) to explore the relationship between reasons and oughts.³ I adopt this model only preliminarily, in order to illustrate some of its limitations. I will introduce the model in this section before exploring its limitations and possible revisions in subsequent sections. The model is only rudimentary, but I believe that it is sufficient to illustrate the different difficulties I will consider.⁴

Before introducing the model, it will be helpful to say something about the language I will be using to describe requirements, and the distinction between requirements and oughts. Traditionally, philosophers have taken ought statements to be the paradigmatic statements of requirement. This assumption is prevalent in formal approaches to modelling the relationship between reasons and deontic conclusions, which follow traditional deontic logic in using the 'ought' operator to express strong deontic necessity (e.g., Horty 2012; Goble 2013). Contemporary linguists and philosophers of language, however, have drawn attention to the distinction between weak necessity operators, like 'ought' and 'should', and strong necessity modals like 'must' and 'have to' (Portner 2009: 79–81; Silk 2015: 298–303; Snedegar 2016: 158–162). Ought statements, at least paradigmatically, do not express requirement. As Bernard Williams (1981: 126) puts it, 'not every *ought* is a *must*'.

The distinction between weak and strong necessity is evident in the following two sentences:

1. I ought to do give my money to charity but it's not as if I have to.
2. #I have to give my money to charity but it's not as if I ought to.

Where (1) is perfectly intelligible, (2) seems incoherent, at least if we interpret 'have to' and 'ought' relative to the same source of reasons or requirements. Similarly, the additional clause in sentence (4), below, appears redundant, where the addition in (3) does not.

3. I ought to do my homework, in fact I have to.
4. #I have to do my homework, in fact I ought to.

3. The model developed in Section 2 is equivalent to the version of Horty's model used by Nair (2016: Appendix 2) to illustrate the relation between reasons and oughts, although I use the operator *Must* rather than *Ought* to express requirement.

4. In brief, the logic ignores: (i) the antecedent conditions of reasons (it lacks any account of reasons as a dyadic relationship of the form 'A is a reason for B'); (ii) exclusionary or undercutting reasons; (iii) defeat of a reason by a set of inconsistent reasons; and (iv) reasons for varying the priority of other reasons. For logics that accommodate these features, see Horty (2012) and Tucker (2018).

As Justin Snedegar (2016: 160) notes, the corpus of linguistic evidence seems to support the conclusion that requirements entail oughts and that oughts entail permissions, but that the reverse entailments do not hold. If John must dance at the ball, then he ought to dance at the ball and he is permitted to do so. It is not the case, however, that oughts entail requirements, nor that permissions entail oughts. Even if John ought to dance at the ball, it is not necessarily the case that he must dance. And though it might be true that John is permitted to dance at the ball, it is not necessarily the case that he ought to dance.

One way of responding to this sort of linguistic evidence is to say that it is beside the point. When many philosophers use ‘ought’, they use it stipulatively to refer to what John Broome (2016) calls the ‘central ought of normativity’, which plays a distinguished role in practical philosophy, and which features in the identification of our rational requirements. But failure to pay attention to salient linguistic distinctions risks creating unnecessary philosophical confusion. Furthermore, the case for observing the difference between oughts and requirements is strengthened by the fact that the distinction in linguistic use tracks a similar set of distinctions in common-sense morality (McNamara 1996; Snedegar 2016: 161–162). Sentence (1), for instance, seems to describe a familiar case of supererogation. Many accounts of supererogation associate it with cases in which I ought to do one thing but am permitted to do another (see, e.g., Raz 1975: 164–165; Heyd 1982 refers to this in terms of the ‘good-ought tie-up’). So long as ‘ought’ is taken to express requirement, however, any model of the relationship between reasons and requirements will face difficulties in accounting for supererogation.

The basic model preserves the distinction between requirements, oughts, and permissions (as do the alternative models I consider below). In order to avoid confusion, I will use the operator *Must* in order to express requirement, and *Ought* to express a weaker sort of deontic necessity. I do not mean to rule out the possibility that common-sense morality is mistaken on these matters. Nevertheless, it would be regrettable if our account of practical reasoning were incompatible with these basic distinctions by stipulation. If proponents of the possibility of supererogation (for example) are making a mistake, it does not seem to me to be a logical or conceptual mistake about the relationship between reasons and requirements.

I assume a basic propositional language \mathcal{L} . Sentences in \mathcal{L} will be denoted using capitalized letters A, B, C , and so on. A *reasoning context* is a pair $\langle \mathcal{R}, < \rangle$ involving a set \mathcal{R} of reasons that obtain in a given context and an order $<$ over these reasons. Reasons are denoted $r: A$, meaning that there is a reason r supporting the proposition expressed by the sentence A . A function \mathcal{C} identifies the content of any reason with a sentence in \mathcal{L} . For the reason $r: A$, $\mathcal{C}(r) = A$. The function can also be applied to any set of reasons, so that $\mathcal{C}(\mathcal{R}) = \{\mathcal{C}(r) : r \in \mathcal{R}\}$.

The turnstile \vdash will represent classical logical consequence. The order $<$ is a strict partial order satisfying both the anti-reflexivity and transitivity axioms.⁵ (Read $r < r'$ as meaning that r' is greater in strength or weight than r).

The model treats reasons as ‘austere’; it does not allow for logical relations between or among reasons (Horty 2012: 41–47). It could be extended to accommodate derivative reasons (cf. Nair 2016), but I will focus in this paper on the relationship between non-derivative reasons and all-things-considered requirements, oughts and permissions. A further limitation of the model is it takes the contents of reasons and the deontic operators that they support to be sentences, rather than actions. For the purposes of discussion we will put to one side the question of whether these resources are sufficient to represent agency.

The model exploits the idea of a maximal consistent subset of consequences of undefeated reasons. We define both the set of undefeated reasons and the concept of a maximal consistent subset of consequences as follows.

Undefeated Reasons: For a reasoning context $\langle \mathcal{R}, < \rangle$ a set of undefeated reasons is defined as $Undefeated_{<}(\mathcal{R}) = \{r \in \mathcal{R} : \text{there is no } r' \in \mathcal{R} \text{ such that (i) } r < r' \text{ and (ii) } \mathcal{C}(r') \vdash \neg \mathcal{C}(r)\}$.

Maximal Consistent Subset: A set of sentences \mathcal{M} is a maximal consistent subset of $\mathcal{C}(\mathcal{R})$ just in case \mathcal{M} is consistent, and there is no other subset \mathcal{M}' such that \mathcal{M}' is consistent and $\mathcal{M} \subset \mathcal{M}' \subseteq \mathcal{C}(\mathcal{R})$.

In order to preserve the desired entailment between requirements, oughts and permissions, we must adopt what has come to be known as the ‘disjunctive’ account of all-things-considered requirements (Horty 2012; Goble 2013). A proposition expressed by a sentence must obtain just in case it is entailed by every maximal consistent subset of conclusions of our undefeated reasons

Must:

$\langle \mathcal{R}, < \rangle \Vdash \text{Must}(X)$ just in case $\mathcal{M} \vdash X$ for every maximal consistent subset \mathcal{M} of $\mathcal{C}(Undefeated_{<}(\mathcal{R}))$.

We can then define these entailment operations over the other logical operators of negation, disjunction, and conjunction in the usual way. Permission is defined as the dual of requirement, as it is in standard deontic logic:

Permission:

$\langle \mathcal{R}, < \rangle \Vdash P(X)$ just in case it is not the case that $\langle \mathcal{R}, < \rangle \Vdash \text{Must}(\neg X)$.

5. That is (i) if $r < r'$ and $r' < r''$ then $r < r''$ and (ii) it is not the case that $r < r'$ and $r' < r$.

The proposition expressed by a sentence is permitted just in case it is not the case that the negation of that sentence is entailed as a requirement.

Against the background of this approach to the relationship between reasons, requirements and permissions, we can extend the basic model to accommodate a relationship between reasons and oughts. Oughts are explained in terms of reasons of undefeated strength. To say that I ought to do something is to say that the sentence expressing the corresponding proposition is entailed by *some* reason or set of reasons of undefeated strength (cf. Chisholm 1974: 125; Raz 1999a: 28–32). If a proposition expressed by a sentence is entailed by some maximal set of undefeated reasons, then it ought to be the case.

Ought:

$\langle \mathcal{R}, < \rangle \sim \text{Ought}(X)$ just in case $\mathcal{M} \vdash X$ for some maximal consistent subset M of $\mathcal{C}(\text{Undefeated}_<(\mathcal{R}))$

In essence, this model combines a so-called ‘disjunctive’ account of requirements with what is known as the ‘conflict’ account of ought statements. These accounts have previously only been considered as rival accounts of ought statements (cf. Horty 2012: ch. 3). Here the accounts are applied to requirements and oughts respectively. The apparent ‘weakness’ of ought is captured by defining it in terms of existential, rather than universal, quantification over maximal consistent subsets of consequences of undefeated reasons.

It is easy to see that this account maintains the desired entailment relationship between requirements, oughts, and permissions.

Observation 1: For any reasoning context $\langle \mathcal{R}, < \rangle$, if $\langle \mathcal{R}, < \rangle \sim \text{Must}(X)$, then $\langle \mathcal{R}, < \rangle \sim \text{Ought}(X)$.

Observation 2: For any reasoning context $\langle \mathcal{R}, < \rangle$, if $\langle \mathcal{R}, < \rangle \sim \text{Ought}(X)$, then $\langle \mathcal{R}, < \rangle \sim P(X)$.

(I offer proofs of these observations, along with all of the observations that follow, in the appendix.) Note that $\text{Ought}(X)$ does not entail $\text{Must}(X)$ in the same fashion, since the model allows for cases where we have two or more equally good but inconsistent sets of reasons. In these cases each set of reasons will entail an ought but not the equivalent requirement. In cases in which there are two or more inconsistent, undefeated sets of reasons, the reasons result in a disjunctive all things considered requirement.

These basic features of the model can be illustrated with any scenario in which an agent is faced with two undefeated reasons, the consequences of which are inconsistent. Consider, for example, Ruth Barcan Marcus’s (1980: 125) ‘Burdan case’ involving the lives of two identical twins in jeopardy, where tragic

circumstances mean I can only save one. Suppose that both twins are drowning at different ends of a lake, so that if I save one I will be abandoning the other. Let T_1 express the proposition that I save the first twin and T_2 the inconsistent proposition that I save the second twin. Our reasoning context $\langle \mathcal{R}, < \rangle_1$ will be such that $\mathcal{R} = \{r_1 : T_1, r_2 : T_2\}$ where the priority order $<$ is empty. Note that $\mathcal{C}(\text{Undeatead}_<(\mathcal{R})) = \{T_1, T_2\}$, with two maximal consistent subsets $\{T_1\}$ and $\{T_2\}$, and that therefore $\mathcal{M} \vdash T_1 \vee T_2$ for both maximal consistent subsets of $\mathcal{C}(\text{Undeatead}_<(\mathcal{R}))$. Therefore $\langle \mathcal{R}, < \rangle_1 \vdash \text{Must}(T_1 \vee T_2)$ and $\langle \mathcal{R}, < \rangle_1 \vdash P(T_1)$ and $\langle \mathcal{R}, < \rangle_1 \vdash P(T_2)$. Since \vdash is reflexive, we also have that $\{T_1\} \vdash T_1$ for the first maximal consistent subset and that $\{T_2\} \vdash T_2$ for the second maximal consistent subset, meaning $\langle \mathcal{R}, < \rangle_1 \vdash \text{Ought}(T_1)$ and $\langle \mathcal{R}, < \rangle_1 \vdash \text{Ought}(T_2)$. I must save one of the two twins. I ought to save the first twin and I ought to save the second twin, and it would be permissible for me to do either. Unfortunately, I cannot do both. Moreover, since requirement and permission are duals, we have that $\langle \mathcal{R}, < \rangle_1 \vdash \neg P(\neg(T_1 \vee T_2))$. It would not be permissible for me save neither twin.

This is, I concede, an attractive picture of the way in which reasons explain permissions, oughts and requirements. It allows us to distinguish our all-things-considered requirements from what we ought to do, and what we ought to do from what is simply permissible, and to do so in a relatively straightforward way. It offers an account of practical reasoning that is consistent with what appears to be the common belief that what we must do is explained by our strongest reasons for action.⁶ It accommodates the distinction between requirements and oughts in a manner that preserves the standard entailments. Nonetheless the model faces difficulties in accommodating various intuitions from common-sense morality about the relationship between reasons, oughts and requirements.

3. Optionality, Tragedy, and Supererogation

The basic model embeds the assumption that the function of all reasons, absent another reason of equal or greater strength, is to generate requirements and therefore to render it impermissible not to conform to them. For any reasoning context $\langle \mathcal{R}, < \rangle$, if there is a reason r such that its conclusion is not inconsistent with the conclusion of any other set of reasons, it will be impermissible not to conform with it. The model therefore reflects the view that reasons, unless defeated or matched by a reason of equal strength, generate requirements (cf. Kagan 1989: 64–70; Schroeder 2007: 130–131). This is, admittedly, a common

6. For an explicit defence of this idea see Portmore (2013).

picture of the function of practical reasons.⁷ But common-sense morality seems to present cases where we have undefeated reasons—even apparently strong or weighty reasons—which it is permissible for us not to act on. Below I discuss several such cases. The cases I consider present challenges that deserve to be considered in greater depth than they are here. My purpose in this section is to set out these problems and observe some possible responses open to those who wish to preserve the basic model. While I am critical of these responses, I do not think that my criticisms are decisive. I simply use them to motivate the further frameworks that are developed in Sections 4 and 5 of the article.

3.1. Optionality

The problem of optionality begins with Raz's 'basic belief' that at any point in time we may permissibly choose between any number of rationally attractive options. As I write this article, it would be permissible for me to do any number of things instead. I could, if I wished, take a break at any moment to have coffee with a colleague. I could begin to prepare for my teaching this afternoon, or I could pause for a moment and plan for my weekend. None of these actions would be contrary to reason, each of them and perhaps several other options are rationally intelligible, and conforming with any one of them would be rationally permissible. If the belief is as basic as Raz asserts, then it places an important constraint on our theorising about reasons: we must explain how it comes to be that, at any point in time, I can permissibly choose between a set of options without being required to perform any one of them. Put another way, some reasons seem to make actions rationally intelligible without making it impermissible not to perform them.

In order to accommodate optionality within the basic model I introduced in Section 2, we need to posit widespread ties or incommensurability in the strength of reasons. The latter approach is favored by Raz (1999b), who takes the basic belief to be a primary motivation for positing widespread incommensurability among reasons. Perhaps the best-known proponent of the former approach is Douglas Portmore (2013), who has developed a sophisticated version of consequentialism that allows for ties in the strength of reasons, once our reasons for action are properly specified. The ability of the basic model to accommodate optionality in terms of ties or incommensurability in the strength of reasons can be illustrated with an example that I borrow from Raz (1999b: 102–103). Mary has an option of going to the theatre this evening, but could also go to visit her mother. Her options are

7. Gert (2004: 19, fn. 3) assembles a list of philosophers who have explicitly or tacitly endorsed this view.

explained by the fact that she has at least two reasons of equal or incommensurable strength which apply to her, neither of which defeats the other. This makes it permissible for her to do one or the other. Raz's example can be illustrated with a simple reasoning context. Let Mary's reasoning context be represented as a structure $\langle \mathcal{R}, < \rangle_2$, where $\mathcal{R} = \{r_1 : M, r_2 : T\}$ where M is inconsistent with T . Suppose that neither her reason to go to the theatre, $r : T$, nor her reason to visit her mother $r : M$, is stronger than the other, so that $<$ is empty. It is easy to see that $\langle \mathcal{R}, < \rangle_2 \models \text{Must}(M \vee T)$, Mary must to go to the theatre or visit her mother, as well as that $\langle \mathcal{R}, < \rangle_2 \models \text{Ought}(M)$ and $\langle \mathcal{R}, < \rangle_2 \models \text{Ought}(T)$. It is, however, not permissible that she neither go to the theatre nor visit her mother $\langle \mathcal{R}, < \rangle_2 \models \neg P(\neg(M \vee T))$, since she must choose one of these two options. (Of course in real life we may face much larger sets of good and inconsistent options, and the basic model will accommodate these cases in terms of a longer disjunctive requirement.)

A great strength of this approach to optionality is that it does not require any departure from the relatively austere basic model introduced above. In fact, theories like those I have attributed to Raz and Portmore can be relatively easily accommodated within the basic model, which therefore allows for a parsimonious reasons-based explanation of optionality. It also seems to accord with a deontic interpretation of the standard decision-theoretic approach: provided that any member of a set of available options is not strictly dominated, it is rationally permissible to choose any one of them, but required to choose between them.

I will have little more to say here about the merits of either Raz or Portmore's proposals for accommodating optionality. The plausibility of their responses to the basic belief rests on substantive arguments, and consideration of both proposals is best left for another occasion. It is worth noting, though, that neither response to the problem of optionality is without philosophical cost. Discomfort with the kind of widespread incommensurability posited by Raz provides one motivation for departing from the view that all reasons play a requiring role (Gert 2004: 102–105). Portmore's arguments for a version of consequentialism that accommodates the basic belief are nuanced and impressive. However, they do rely on assumptions about the relationship between rational and moral reasons that are somewhat controversial (cf. Portmore 2013: 440–441) and about the plausibility of assessing our options indirectly, in respect of the different maximal options they comprise (cf. Gert 2014: 215–216).

3.2. Distinguishing Optionality from Tragic Conflicts

Even if we accept the basic model's account of optionality, we will face the further problem of distinguishing cases of optionality from cases of apparently

tragic conflict, like the drowning twins scenario considered in Section 2. Recall that in the drowning twins scenario the reasoning context supported the conclusion that rescuing one of the two twins was required: I ought to rescue the first twin and I ought to rescue the second twin, but failing to rescue one of them is permissible. Within the basic model, the deontic conclusions generated by the drowning twins scenario are structurally identical to Mary's choice between visiting her mother and going to the theatre. Mary must either visit her mother or go to the theatre, ought to go to the theatre and ought to visit her mother, but may permissibly fail to perform one of these options. Conflict between reasons that would otherwise generate requirements lends them a kind of optionality: conflicts like the drowning twins scenario are no different from a situation like Mary's where she was forced to choose between two good, undefeated options. Both cases involve situations in which we are able to conform to only one of two or more inconsistent reasons and may permissibly choose between them. But the drowning twins scenario seems different in a normatively significant way. We are not merely choosing between two goods, we are forced to make a tragic choice between two equally unattractive options (cf. Dancy 1993: 123). It appears as though an account of practical reasoning should be able to accommodate the difference between these two contexts. The basic model's method of distinguishing between oughts and requirements does not allow us to do so.

One way of accommodating tragic conflicts would be to depart from the assumption that requirements entail permissions in favor of an account that allows for conflicting all-things-considered requirements (Bedke 2011: 149–151; Horty 2012: ch. 4). The tragedy of these cases could then be explained by their involving a choice between two impermissible options. I am sympathetic to the argument that any reasons-based account should be able to accommodate this sort of unresolved conflict between all-things-considered requirements. For the purposes of discussion, however, I will not explore this possibility, because I am interested in developing a model for practical reasoning that can preserve the standard inference relationship between reasons, oughts and permissions that I discussed in Section 2. The intuitions underlying these entailments seem to be widely accepted, and it would be problematic if our framework of practical reasoning could not accommodate them. In any case, the problem of distinguishing between optionality and tragic conflicts will still arise for proponents of the possibility of conflicting all things considered requirements. Where the basic model treats tragic conflicts as cases of optionality, a version of the same model that allows for conflicting requirements will treat all cases of optionality as involving a tragic choice between two impermissible options.

3.3. Supererogation and Justification

Various forms of supererogation also suggest the possibility of reasons that lack requiring force. Consider the situation of someone who decides to sacrifice their own life to save the lives of several others. Many people share the judgement that the person who sacrifices their life in this way has acted justifiably, in a way that is rationally permissible. We are, however, reluctant to classify this sort of self-sacrifice as the course of action that we *must* take. Within the basic model, if a reason is stronger than another inconsistent reason, the reason supports a statement of requirement. Suppose that I face a choice between saving the lives of several strangers by throwing myself on a grenade or simply running for safety in an attempt to save my own life. Let G express the proposition that I throw myself on the grenade, and R the inconsistent proposition that I run for safety. Now suppose that we have a reasoning context $\langle \mathcal{R}, < \rangle_3$ where $\mathcal{R} = \{r_1 : R, r_2 : G\}$, where $R \vdash \neg G$ and where $r_1 < r_2$, so that my reason to save the lives of others is stronger than my reason to run for safety. It is easy to see that in this context, $\langle \mathcal{R}, < \rangle_3 \vdash \text{Must}(G)$ and that $\langle \mathcal{R}, < \rangle_3 \vdash \neg P(R)$. The reasoning context supports the conclusion that I must save the lives of the strangers, and that it is impermissible for me to run for safety.

Cases of supererogation can be accommodated within the basic model by positing that they involve a conflict between different classes of reasons. Portmore (2013) suggests that supererogation arises because we can have greatest moral reason to act in a way that is inconsistent with our greatest non-moral reason. Such a response relies on the availability of a robust distinction between moral reasons and non-moral reasons. It also has difficulty with cases in which apparently weaker moral reasons appear to make it permissible not to act on a stronger reason (Gert 2014: 212). For instance, suppose I decide to give some proportion of my money to a moderately effective local charity rather than using it to fund very effective global malaria prevention efforts. Although my reason to give to the local charity is a weaker moral reason than my reason to fund malaria prevention, it still seems that I am justified in giving to the local charity.

Justifications for action that would normally be impermissible provide another, less dramatic example of a structurally similar phenomenon. Even when they have great strength, the reasons associated with justifications do not require us to act in accordance with the justification. If I strike someone in self-defence then I have acted justifiably, but it seems implausible to suggest that it would have been impermissible for me not to strike them. A reasons-based account of justifications ought to be able to distinguish what is merely permissible from what ought to occur, and what ought to occur from what is required. But, on the basic model, the only way to incorporate a justification with non-requiring

strength is to postulate that the justifying reason is of equal or incomparable strength to another reason.

Following Raz (1975; 1999a: 90–96), we might seek to amend the basic model by appealing to the possibility of ‘exclusionary permissions’, which give us reasons not to act on otherwise good reasons, in order to accommodate supererogation and justification within the basic framework. My reason to strike someone in self-defence could be paired with an exclusionary permission not to act on my reason not to strike them. Similarly, my reason to act out of concern for my own life could be paired with a permission not to act on my reason to jump on the grenade (for a formal account of exclusionary reasons see Horty 2012: 122–130). One problem with this form of explanation is that if my reason not to strike the assailant is excluded, we now lack an explanation of why it would nonetheless be permissible for me to fail to strike them, notwithstanding my remaining good reason to act in self-defence. If I act on my reason not to strike in retaliation, I have acted on a reason that I had good reason to exclude from deliberation. If the reason to exclude is unmatched or undefeated, then it would appear to require me to exclude my good reason not to strike my assailant. If I am not required to exclude this reason (cf. Raz 1999a: 90), then the proponent of exclusionary permissions must offer an explanation of the optional character of the reason to exclude. The problem of optionality reappears. In the case of supererogation, moreover, an appeal to exclusionary permissions raises formal difficulties about the relation between an exclusionary reason and the strength of the reasons that it excludes. How is it that my reason to sacrifice myself for others, although stronger than my reason to run for safety, comes to be excluded? One response would involve appealing to the problematic idea that weaker reasons for action can exclude or undercut weightier reasons (cf. Horty 2012: 135–141; Tucker 2018: 959–962). Another would be to posit that the exclusionary reason is weightier than my reason to sacrifice myself. The proponent of exclusionary permissions is then left with the problem of justifying the weight of this exclusionary reason. A proponent of exclusionary permissions is unlikely to regard these difficulties as insurmountable, but they suggest to me that alternative formal approaches are worth considering.

4. Appealing to Reasons’ Strength

A seemingly plausible response to these problems is to appeal to the strength of reasons in order to distinguish between reasons that possess justifying force and reasons that possess requiring force. It is tempting to suggest that some reasons are insufficiently strong to support requirements, and that while weaker reasons

support oughts, only reasons that possess strength above a certain threshold will support requirements.⁸

Scanlon (1998; 2014), Schroeder (2007), and Skorupski (2010) have all developed theories according to which reasons possess degrees of weight or strength, so appealing to a threshold degree of strength in this way seems plausible. In his more recent work, Scanlon (2014: 111, fn. 8) seems to develop an account of the distinction between requiring and non-requiring reasons by appealing to their relative weight. Scanlon follows Gert in distinguishing between reasons with justifying strength and reasons with requiring strength but argues that these are not different ‘ways in which a consideration can count in favor of an action’ (2014: 111 fn. 8). According to Scanlon, the difference between requiring and non-requiring reasons is that requiring reasons are sufficiently strong that few reasons could justify my failure to conform to them.

The idea of appealing to strength of reasons to accommodate latitude can be sketched with a slight alteration of the model introduced above. Supposing that reasons are partially ordered in terms of priority or ‘weight’, we can then postulate a threshold value, above which the reasons in question have the quality of a deontic requirement, and below which conformity is optional. Reasons below threshold can support oughts, but not requirements.

Let a *threshold reasoning context* be a structure $\langle \mathcal{R}, <, n \rangle$, where $<$ is a strict partial order over reasons, as above, except that it now contains the privileged value n , which acts as a threshold for the assignment of requiring force to any reason. Let \mathcal{R}_n be the set of requiring reasons in \mathcal{R} above this threshold. We can then adopt an account of requirements that refers only to the subset of requiring reasons applicable in a given scenario, while still allowing for reasons below threshold to make actions rationally attractive or intelligible, and therefore to support ought statements.

More formally, we define a set of requiring reasons for a threshold reasoning context as follows:

Requiring Reasons: For a threshold reasoning context $\langle \mathcal{R}, <, n \rangle$, a set of requiring reasons \mathcal{R}_n is defined as $\mathcal{R}_n = \{r \in \mathcal{R} : n < r\}$.

We then define our requirement operator as for the basic model, this time restricting our definition so that it quantifies only over those reasons that are above threshold.

8. Both Gert (2004: 92–101) and Snedegar (2016: 162–163) also suggest, and then reject, proposals of this sort.

Must (Threshold Reasoning Context): $\langle \mathcal{R}, <, n \rangle \vdash \text{Must}(X)$ just in case $\mathcal{M} \vdash X$ for every maximal consistent subset \mathcal{M} of $\mathcal{C}(\text{Undeated}_{<}(\mathcal{R}_n))$.

The permission operator P is then defined as the dual of *Must* as before.

In order to accommodate oughts within this approach, I will stipulate that an act ought to be performed just in case it is supported by a maximal set of undefeated reasons that is consistent with one of our maximal consistent sets of requiring reasons. (Otherwise the account will allow for ought statements that are inconsistent with our requirements). In order to capture this idea in our definition of oughts, I will appeal to the notion of a maximal *requirement consistent* subset of undefeated reasons. A set of reasons will be requirement consistent just in case there is some maximal consistent subset of requiring reasons that is consistent with it.

Requirement Consistency (Threshold Reasoning Contexts): Where $\langle \mathcal{R}, <, n \rangle$ is a threshold reasoning context, a set of sentences \mathcal{M} is requirement consistent just in case $\mathcal{M} \cup \mathcal{M}'$ is consistent for some maximal consistent subset \mathcal{M}' of $\mathcal{C}(\mathcal{R}_n)$.

We then stipulate that maximal requirement-consistent subsets of consequences of undefeated reasons support ought statements.

Maximal Requirement Consistent Subset: A subset \mathcal{M} of a set of consequences of reasons $\mathcal{C}(\mathcal{R})$ is a maximal requirement consistent subset of $\mathcal{C}(\mathcal{R})$ just in case (i) \mathcal{M} is requirement consistent and (ii) there is no other subset \mathcal{M}' such that \mathcal{M}' is requirement consistent and $\mathcal{M} \subset \mathcal{M}' \subseteq \mathcal{C}(\mathcal{R})$.

Ought (Threshold Reasoning Context): $\langle \mathcal{R}, <, n \rangle \vdash \text{Ought}(X)$ just in case $\mathcal{M} \vdash X$ for some maximal requirement-consistent subset \mathcal{M} of $\mathcal{C}(\text{Undeated}_{<}(\mathcal{R}))$.

Defining oughts in terms of requirement consistency ensures that the contents of an ought statement are never inconsistent with the contents of a requirement supported by the same reasoning context.

Observation 3: Where $\langle \mathcal{R}, <, n \rangle$ is any threshold reasoning context, if $\langle \mathcal{R}, <, n \rangle \vdash \text{Must}(X)$ and $\langle \mathcal{R}, <, n \rangle \vdash \text{Ought}(Y)$, then it is not the case that $X \vdash \neg Y$.

Appealing to requirement consistency in this way also ensures that threshold reasoning contexts preserve a number of important patterns of reasoning that result from the agglomeration of oughts with requirements. More formally, if a reasoning context supports the conclusion $\text{Must}(X)$ and $\text{Ought}(Y)$, then it also

supports the conclusion $Ought(X \wedge Y)$. This allows threshold reasoning contexts to preserve some attractive patterns of inference. I discuss this aspect of the model further in the appendix.

It is also easy to see that this account preserves the same entailments as the basic model. If a threshold reasoning context supports a requirement then it supports a corresponding ought statement, but not vice versa.

Observation 4: For any threshold reasoning context $\langle \mathcal{R}, <, n \rangle$, if $\langle \mathcal{R}, <, n \rangle \sim Must(X)$, then $\langle \mathcal{R}, <, n \rangle \sim Ought(X)$.

Observation 5: For any threshold reasoning context $\langle \mathcal{R}, <, n \rangle$, if $\langle \mathcal{R}, <, n \rangle \sim Ought(X)$, then $\langle \mathcal{R}, <, n \rangle \sim P(X)$.

The model allows us to distinguish cases of optionality from cases of tragic conflict by appealing to reasons of non-requiring strength. Suppose that Mary's scenario, where she has good reason to go to the theatre or to visit her mother, is now represented by $\langle \mathcal{R}, <, n \rangle_1$ a threshold reasoning context such that $\mathcal{R} = \{r_1 : M, r_2 : T\}$ where both $r_1 < n$ and $r_2 < n$, so that Mary's reason to see her mother and her reason to go to the theatre are below threshold. The context $\langle \mathcal{R}, <, n \rangle_1$ supports the conclusions $Ought(M)$ and $Ought(T)$, but not the conclusion $Must(M \vee T)$. Therefore, it supports the conclusions $P(\neg M)$ and $P(\neg T)$. Mary ought to visit her mother and ought to go the theatre, but it is permissible for Mary to do neither.

Appealing to threshold strengths seems to hold some promise for accommodating optionality, in particular. When used to augment the basic model in the manner suggested, it allows us to distinguish tragic cases like the drowning twins scenario from cases in which someone must merely choose between two good options, and so it allows us to remedy one of the principal drawbacks for the basic model that I outlined above.

However, I doubt that the difference between the requiring and recommending force of reasons can be captured through an appeal to strength alone. To begin with, whether or not reasons play a requiring or justifying role does not seem to track their apparent strength. It seems plausible that some reasons can possess a great deal of strength without acquiring deontic or requiring force. Suppose that I have nothing to do this weekend, and that going to the cinema or to the beach may bring me a great deal of pleasure, far more than if I choose to stay at home. I think it would nonetheless not be impermissible or irrational for me to fail to go to the beach or cinema. My reasons to go out to the beach or cinema are strong, but they do not seem to require conformity.

The distinction between the strength of a reason and its requiring force is apparent when we consider cases of supererogation (Gert 2004: 88–92; Snedegar 2016: 163). Recall the example of the person who has the option of sacrificing

their own life in order to save the lives of many others. Their reasons for self-sacrifice in these situations seem particularly strong—far stronger, certainly, than their reason to preserve their own life. Someone who did sacrifice their own life in this way would certainly have acted justifiably. Regardless of the strength of this reason, however, it does not seem to require self-sacrifice. Proponents of the threshold approach need to give some account of why these apparently stronger reasons fail to acquire requiring strength, even though they apparently defeat or match other reasons of requiring strength. Cases in which reasons justify what would normally be impermissible seem to have a similar structure. For instance, my reason to act in self-defence seems to have far greater strength than my reason to allow someone else to strike me, but in spite of its strength it does not make it impermissible for me to allow someone else to strike me. The person who turns the other cheek does not seem to make any rational or moral error. In spite of its apparent strength, our reason to act in self-defence does not acquire any requiring force.

I do not mean to dismiss the prospects of a strength-based approach to the distinction between oughts and requirements outright. I have only sketched a very crude version of such a model here, and it may well be that a more plausible version of the model can be developed. However, to the extent that we want to accommodate the possibility of justifications and supererogation within reasons-based account, these cases give us reason to doubt that appealing to reasons' strength in order to distinguish between reasons' justifying and requiring roles will be sufficient. Instead, I will propose that we modify our models in order to accommodate the fact that reasons can have differing degrees of justifying and requiring strength.

5. Distinguishing Reasons' Justifying and Requiring Roles

Various authors have been attracted to the idea that it is necessary to postulate a class of reasons that possess non-requiring force. For instance, Jonathan Dancy (2004), Ulrike Heuer (2004), Margaret Little and Colleen Macnamara (Little 2013; Macnamara & Little 2017) have suggested that it is necessary to postulate a separate class of 'enticing', 'optional', or 'commendatory' reasons that recommend an action without making it impermissible not to pursue that action. Joshua Gert (2004; 2007) and Patricia Greenspan (2005; 2007; 2010) make the related suggestion that some reasons function to justify actions without requiring conformity. As Gert puts it, some reasons are 'purely justifying reasons', in that they have a strong justificatory function without having any requiring force (2004: 23). The model I develop in this section is based on Gert and Greenspan's insights, though it is not intended to be entirely faithful to either author's views.

As Gert notes (2016: 158), it is philosophically uncontroversial that reasons can occupy both justifying and requiring roles. Advocates of the basic model do not deny this. Accounts that appeal to ‘favoring’ or ‘justifying’ reasons differ from the basic model in allowing that reasons can occupy a justifying role without possessing any requiring force (cf. Gert 2007: 534–537; Greenspan 2010: 190). Unlike the strength-based approaches discussed in the previous section, advocates of what I will call ‘dual-role’ approaches allow that reasons’ strength need not correspond to any requiring role in a given context. Even strong reasons can play a justificatory role without playing any requiring role. For instance, suppose that I decide to break my promise to meet someone for lunch because I am unwell. My reason to keep my promise plays a requiring role—it ordinarily supports a requirement to perform the promised act. The fact that I am feeling unwell, on the other hand, gives me a purely justificatory reason of sufficient strength not to comply with my promise and to stay at home and rest. My reason to stay at home and convalesce defeats my reason to keep my promise. Nonetheless, my reason to stay at home does not play any requiring role. It would be rationally permissible for me to stay at home and convalesce, but it is not the case that I *must* stay at home. On the other hand, where a requiring reason outweighs or defeat a reason with only justifying strength it can make what would otherwise be a rationally justified option rationally impermissible. If drinking a cup of coffee would make me feel unwell, my reason not to drink the coffee has sufficient strength to defeat my reason to drink the coffee, even though my reason to drink the coffee would normally render it permissible for me to do so (Gert 2007: 537).

These insights can be used to fashion an extension of the basic model that allocates to each reason a justifying or requiring role (or both). In order to accommodate both dimensions of strength, we will postulate *dual-role reasoning contexts* $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle$. The two sets \mathcal{D} and \mathcal{J} identify all reasons $r \in \mathcal{R}$ that occupy requiring and justifying roles, respectively. The priority order $<$ is, as for our basic reasoning contexts, a strict partial order reflecting reasons’ weight or strength, regardless of the roles that a reason occupies.⁹

Within these models, requirements are generated by reasons of undefeated strength that occupy the requiring role. Only those reasons that play a requiring role and have undefeated strength support requirements. I will describe such reasons as *undefeated requiring reasons*.

9. Gert sometimes invokes the idea that that the same reason can possess different ‘degrees’ of justifying and requiring strength, where one reason can possess more justifying strength than another reason possesses requiring strength (e.g., 2007: 546–548). This suggests the possibility of cross-scalar comparison. On the account I offer here, however, a reason’s strength carries across into whatever role it occupies. Like Greenspan (2005: 390, fn. 1), I regard the characterization of a reason’s strength as entirely separate from its occupation of either a justifying or requiring role.

Undeclared Requiring Reasons: For a dual-role reasoning context $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle$, a set of undeclared requiring reasons is defined as $Undeclared_{<}(\mathcal{D}) = \{r \in \mathcal{D} : \text{there is no } r' \in \mathcal{R} \text{ such that (i) } r < r' \text{ and (ii) } \mathcal{C}(r') \vdash \neg \mathcal{C}(r)\}$.

Essentially, a requiring reason is undeclared just in case there is no other inconsistent reason in \mathcal{R} of greater priority (regardless of whether that other reason also occupies the requiring role). A proposition expressed by a sentence will be required if it is supported by some maximal consistent subset of reasons with undeclared requiring strength.

Must (Dual-Role Reasoning Contexts): $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \sim Must(X)$ just in case $\mathcal{M} \vdash X$ for every maximal consistent subset \mathcal{M} of $\mathcal{C}(Undeclared_{<}(\mathcal{D}))$.

Where the permission operator P is again defined as the dual of $Must$.

On this approach, a reason can occupy only a justifying role and nonetheless defeat a requiring reason when it has greater strength. The purely justificatory reason will not support a requirement, but will still be sufficient to defeat a weaker requiring reason. In its current form, however, this model lacks any account of the relationship between reasons and oughts. For this reason both Matt Bedke (2011: 134) and Snedegar (2016: 169–172) criticise Gert’s approach for lacking an account of the relationship between reasons and what we ought to do. The model certainly requires us to depart from the simple idea that we ought to do what we have most reason to do. It does, however, allow us to offer an account of ought in terms of our best justifying reasons. On the particular interpretation I favor, we ought to conform to one of our strongest compatible sets of justifying reasons (cf. Gert 2014: 209). In other words, a dual-role reasoning context will support an ought just in case it is supported by some maximal consistent subset of undeclared justifying reasons. In order to develop this account of oughts, we first define a set of undeclared justifying reasons.

Undeclared Justifying Reasons: For a dual-role reasoning context $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle$ a set of undeclared justifying reasons is defined as $Undeclared_{<}(\mathcal{J}) = \{r \in \mathcal{J} : \text{there is no } r' \in \mathcal{R} \text{ such that (i) } r < r' \text{ and (ii) } \mathcal{C}(r') \vdash \neg \mathcal{C}(r)\}$.

As with our threshold reasoning contexts, we also need to stipulate that any maximal set of justifying reasons that supports an ought will need to be requirement consistent, otherwise the standard entailments will not be preserved.¹⁰

10. Consider the model $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle_{\leq}$, where $\mathcal{D} = \{r_1 : A, r_2 : B\}$ and $\mathcal{J} = \{r_1 : A, r_2 : B, r_3 : C\}$, where $<$ is empty, and where $A \vdash \neg B$ and where $B \wedge C \vdash D$. We have that $\mathcal{C}(Undeclared_{<}(\mathcal{D})) = \{A, B\}$

To do so, we first alter the definition of requirement consistency for a dual-role reasoning context:

Requirement Consistency (Dual-Role Reasoning Context): Where $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle$ is a dual-role reasoning context, a set of sentences \mathcal{M} is requirement consistent just in case $\mathcal{M} \cup \mathcal{M}'$ is consistent for some consistent subset \mathcal{M}' of $\mathcal{C}(\text{Undeatead}_{<}(\mathcal{D}))$.

We then define oughts in terms of maximal requirement consistent subsets, as before:

Ought (Dual-Role Reasoning Contexts): $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \sim \text{Ought}(X)$ just in case $\mathcal{M} \vdash X$ for some maximal requirement consistent subset \mathcal{M} of $\mathcal{C}(\text{Undeatead}_{<}(\mathcal{J}))$.

On this account of our practical reasoning, we first identify our undefeated requiring reasons, in order to determine what is required or impermissible. Oughts are then supported by our best justifying reasons, provided the consequences of their conclusions are consistent with some maximal subset of requiring reasons. As for our threshold reasoning contexts, invoking requirement consistency prevents our models from generating requirements and oughts with inconsistent contents.

Observation 6: Where $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle$ is any dual-role reasoning context, if $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \sim \text{Must}(X)$ and $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \sim \text{Ought}(Y)$, then it is not the case that $X \vdash \neg Y$.

As I note in the appendix, introduction of requirement consistency also allows for the agglomeration of oughts with requirements (just as it did for the threshold reasoning contexts), thereby preserving attractive patterns of inference.

Dual-role reasoning contexts are able to accommodate the problematic cases we noted for the basic model in Section 3, as I demonstrate below. However, as Snedegar (2016: 170–171) notes in his discussion of Gert's approach, unamended it is unable to accommodate the standard entailment between requirements and oughts. Since a reason could occupy the requiring role without occupying a justifying role, it follows that a reason could fail to belong to the set of undefeated justifying reasons, and therefore fail to support an ought. Consider, for example,

and $\mathcal{C}(\text{Undeatead}_{<}(\mathcal{J})) = \{A, B, C\}$. In this dual-role context we have that $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle_5 \sim \text{Must}(\neg D)$ but that $\mathcal{M} \vdash D$ for some maximal consistent subset \mathcal{M} of $\text{Undeatead}_{<}(\mathcal{J})$, and therefore that $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle_5 \sim \text{Ought}(D)$.

a dual-role reasoning context $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle_1$, an interpretation of the drowning twins scenario, where $\mathcal{D} = \{r_1 : T_1, r_2 : T_2\}$ but where $\mathcal{J} = \emptyset$. The reason to save either twin occupies a requiring role but not a justifying one. It is easy to see that this context will generate the conclusion $Must(T_1 \vee T_2)$ but not the conclusion $Ought(T_1 \vee T_2)$, since neither r_1 nor r_2 are in $Undeclared(\mathcal{J})$.

Fortunately it can be shown that, granted a plausible constraint on their structure, dual-role reasoning contexts preserve the desired entailments between requirements, oughts and permissions. Provided we impose the constraint on our contexts that any reason that occupies a requiring role also occupies a justifying role, then any set of reasons supporting the conclusion $Must(X)$ will also support the conclusion $Ought(X)$. More formally, we stipulate that for any reasoning context $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle, \mathcal{D} \subseteq \mathcal{J}$. This stipulation seems plausible, since it would be odd for a reason that occupies the requiring role not to also occupy the justifying role (Greenspan 2010: 189).¹¹ Any reason that can require us to perform an action can also provide us with justification for the same action.

Observation 7: Suppose $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle$ is any dual-role reasoning context such that $\mathcal{D} \subseteq \mathcal{J}$. If $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash Must(X)$, then $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash Ought(X)$.

Moreover, once this stipulation is made, any set of reasons that supports the conclusion $Ought(X)$ will also support the conclusion $P(X)$.

Observation 8: Where $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle$ is any dual-role reasoning context, if $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash Ought(X)$, then $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash P(X)$.

Unsurprisingly given the motivations for developing the model, dual-role reasoning contexts generate plausible conclusions in the cases of optionality, supererogation and justification that presented problems for the basic model. Latitude emerges whenever there are undefeated reasons that do not occupy a requiring role (Greenspan 2010: 188–189). Consider the case of Mary, introduced above, in which it appears as though she has two good, non-requiring options to go to the theatre or to see her mother. Let our reasoning context be $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle_2$, where as before $\mathcal{R} = \{r_1 : M, r_2 : T\}$, but where $\mathcal{D} = \emptyset$, and where as before T is consistent with M . It is easy to see that $\mathcal{C}(Undeclared(\mathcal{J})) = \{M, T\}$ but that $\mathcal{C}(Undeclared(\mathcal{D})) = \emptyset$. Therefore we have that $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle_2 \vdash Ought(T)$ and $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle_2 \vdash Ought(M)$ as

11. As a reviewer notes, this stipulation makes the separate set \mathcal{J} of justifying reasons redundant. The same results could be achieved with a model $\langle \mathcal{R}, \mathcal{D}, < \rangle$, which distinguished between ordinary reasons and requiring reasons. However, I retain a distinct set of justifying reasons for the purposes of exposition.

well as that $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle_2 \vdash \sim P(\neg(T \vee M))$. Mary ought to go to the theatre and ought to see her mother, but it is permissible for her to do neither.

Cases of supererogation similarly arise in contexts where the supererogatory options occupy a justificatory role and have greater strength than other options, but where they do not have any requiring force and therefore do not support a requirement. For instance, consider a dual-role reasoning context $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle_3$ where $\mathcal{J} = \{r_1 : G, r_2 : R\}$, where as before G expresses the proposition that I save the lives of several strangers by jumping on a grenade, and where sentence R expresses the inconsistent proposition that I run away. The priority order is such that $r_2 < r_1$, my reason to run away is weaker than my reason to save others' lives. In this reasoning context we have $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle_3 \vdash \text{Ought}(G)$, but not $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle_3 \vdash \text{Must}(G)$. It therefore supports the further conclusion $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle_3 \vdash P(R)$. I ought to jump on the grenade to save the lives of several strangers, since that would be consistent with my best justifying reasons, but it is permissible for me to attempt to save myself by running away.

Finally, the revised model provides the resources to distinguish cases of optionality from cases of tragic conflict. In dual-role reasoning contexts, tragic conflict arises when there are two or more conflicting and undefeated requiring reasons, as opposed to two or more conflicting purely justificatory reasons. We can interpret the drowning twins scenario in terms of a dual-role context $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle_4$ where $\mathcal{D} = \{r_1 : T_1, r_2 : T_2\}$, as before, and where $<$ is again empty. Since both reasons are undefeated in the requiring role, we have that $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle_4 \vdash \text{Must}(T_1 \vee T_2)$ and therefore that $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle_4 \vdash \neg P(\neg(T_1 \vee T_2))$. This is to be contrasted with a context, like Mary's case considered above and represented in the dual-role reasoning context $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle_2$, in which we merely choose between two good options which are undefeated justifying reasons, but which do not function as requiring reasons.

I have presented this model as an interpretation of Gert and Greenspan's suggestion that some reasons perform a justifying role without having any requiring force. I do not mean to suggest that either Gert or Greenspan would endorse the model as I have developed it. For instance, the definition of oughts I have adopted allows the model to accommodate the idea that we ought to conform to one of our best justifying reasons. This aspect of the model allows us to offer an account of the relationship between justifying reasons and oughts. Gert would likely reject this particular aspect of the model, since he rejects the idea that what we ought to do is in any way associated with our strongest reasons as misconceived (2007: 548–552).¹² Rather than attempting to capture Gert and

12. In more recent work, Gert (2012: 616–617) seems sympathetic to the view that, in addition to occupying requiring and justifying roles, reasons might also perform a merit-conferring role (this view is defended in Horgan & Timmons 2010).

Greenspan's particular frameworks, I have provided a more generic approach that exploits their insight that reasons perform either or both justifying and requiring roles in a given context. As we have seen, the model can accommodate the cases that are most problematic for the basic model.

6. Conclusion

My focus in this article has been on developing simple reasons-based models that accommodate the various problems from common-sense morality that I noted in Section 3, while also preserving the desired entailment relationship between requirements, oughts and permissions. Focussing on these issues has constrained the scope of my discussion in a number of ways. The models are only rudimentary. There are plausible extensions of the threshold and dual-role models that I have not considered. Both models could be amended to accommodate familiar forms of reasoning about the priorities or roles that are assigned to other reasons in a given context (cf. Horty 2012; Tucker 2018). For instance, whether a given reason occupies a requiring role in a given context could itself be the subject of further reasoning that is represented within the model. Similarly, while there are plausible variants of the models that allow for conflicting all-things-considered requirements, I have not discussed them.

I have also not discussed some important alternative approaches to addressing the problems I identified in Section 3. Two forms of alternative response seem to me to be particularly important from a reasons-based perspective. First, Bedke (2011) proposes a 'Millian inversion' by which requirements and permissions are explained in terms of what we have most reason to respond to with speech acts or actions requiring or permitting certain behaviour. However, since this response requires us to give up on the standard entailment relationships (Snedegar 2016, 165–167), I have not explored its merits here. Second there is the possibility, which I only briefly critiqued in Section 3, of using Raz's (1975; 1999a) account of exclusionary permissions to account for some of the cases that presented problems for the basic model. I do not think that invoking exclusionary reasons provides a straightforward solution to these problems, but the approach is certainly deserving of further consideration. The relation between exclusionary reasons and the distinction between justifying and requiring reasons merits lengthier discussion, and I hope to revisit it in future work.

There are philosophical problems with the dual-role approach to reasons that cannot be addressed by any formal model. Most obviously, it complicates the simple and attractive idea that the function of a practical reason is to 'count

in favour' of an action (Scanlon 1998: 17). I suspect, for this reason, that many proponents of the basic model will be unmoved by my motivations for developing an alternative account. Nonetheless, the models developed here allow us to identify some points of disagreement between proponents of the basic model and its critics with greater precision. More significantly, they show that a coherent model of the dual-role approach to practical reasoning is available—one that is consistent with a conventional understanding of the relationship between reasons, oughts and requirements.

Appendix: Properties of the Models

A1. The Basic Model

Observation 1: Where $\langle \mathcal{R}, \prec \rangle$ is any reasoning context, if $\langle \mathcal{R}, \prec \rangle \vdash \text{Must}(X)$, then $\langle \mathcal{R}, \prec \rangle \vdash \text{Ought}(X)$.

Proof: If $\langle \mathcal{R}, \prec \rangle \vdash \text{Must}(X)$ then $\mathcal{M} \vdash X$ for all maximal consistent subsets \mathcal{M} of $\mathcal{C}(\text{Undeclared}_{\prec}(\mathcal{R}))$. So clearly $\mathcal{M}' \vdash X$ for some maximal consistent subset \mathcal{M}' of $\mathcal{C}(\text{Undeclared}_{\prec}(\mathcal{R}))$. Thus $\langle \mathcal{R}, \prec \rangle \vdash \text{Ought}(X)$.

Observation 2: Where $\langle \mathcal{R}, \prec \rangle$ is any reasoning context, if $\langle \mathcal{R}, \prec \rangle \vdash \text{Ought}(X)$, then $\langle \mathcal{R}, \prec \rangle \vdash P(X)$.

Proof: If $\langle \mathcal{R}, \prec \rangle \vdash \text{Ought}(X)$ then $\mathcal{M} \vdash X$ for some maximal consistent subset \mathcal{M} of $\mathcal{C}(\text{Undeclared}_{\prec}(\mathcal{R}))$. If $\mathcal{M} \vdash X$ then $\mathcal{M} \not\vdash \neg X$. Thus it is not the case that $\langle \mathcal{R}, \prec \rangle \vdash \text{Must}(\neg X)$. Thus $\langle \mathcal{R}, \prec \rangle \vdash P(X)$.

A2. Threshold Reasoning Contexts

Several proofs below invoke the following intermediate observations, which I prove separately.

Observation*: Where $\langle \mathcal{R}, \prec, n \rangle$ is any threshold reasoning context, every maximal consistent subset \mathcal{M} of $\mathcal{C}(\text{Undeclared}_{\prec}(\mathcal{R}_n))$ is a requirement consistent subset of $\mathcal{C}(\text{Undeclared}_{\prec}(\mathcal{R}))$.

Proof: Suppose $r \in \text{Undeclared}_{\prec}(\mathcal{R}_n)$. For any other $r' \in \mathcal{R}$, either $n < r'$ or not. If $n < r'$, then, since $r \in \text{Undeclared}_{\prec}(\mathcal{R}_n)$, by definition it is either not the case that $r < r'$ or not the case that $\mathcal{C}(r') \vdash \neg \mathcal{C}(r)$. If it is not the case that $n < r'$, then since $r \in \text{Undeclared}_{\prec}(\mathcal{R}_n)$ it cannot be the case that $r < r'$, since $n < r$ and by transitivity we would have that $n < r'$, contradicting our assumption. Thus in either case for any $r \in \text{Undeclared}_{\prec}(\mathcal{R}_n)$ there is no $r' \in \mathcal{R}$ such that both (i) $r < r'$ and (ii) $\mathcal{C}(r') \vdash \neg \mathcal{C}(r)$. Therefore for any $r \in \text{Undeclared}_{\prec}(\mathcal{R}_n)$, $r \in \text{Undeclared}_{\prec}(\mathcal{R})$. Therefore $\text{Undeclared}_{\prec}(\mathcal{R}_n) \subseteq \text{Undeclared}_{\prec}(\mathcal{R})$ and $\mathcal{C}(\text{Undeclared}_{\prec}(\mathcal{R}_n)) \subseteq \mathcal{C}(\text{Undeclared}_{\prec}(\mathcal{R}))$. Thus for any maximal consistent subset $\mathcal{M} \subseteq$ of $\mathcal{C}(\text{Undeclared}_{\prec}(\mathcal{R}_n))$, it is also the case that $\mathcal{M} \subseteq \mathcal{C}(\text{Undeclared}_{\prec}(\mathcal{R}))$. Clearly every maximal consistent subset \mathcal{M} of $\mathcal{C}(\text{Undeclared}_{\prec}(\mathcal{R}_n))$ is also requirement consistent. Therefore every maximal consistent subset \mathcal{M} of $\mathcal{C}(\text{Undeclared}_{\prec}(\mathcal{R}_n))$ is a requirement consistent subset of $\mathcal{C}(\text{Undeclared}_{\prec}(\mathcal{R}))$.

Observation:** Where $\langle \mathcal{R}, <, n \rangle$ is any threshold reasoning context, for every maximal consistent subset \mathcal{M}' of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}))$ there is some maximal requirement consistent subset \mathcal{M} of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}_n))$ such that $\mathcal{M} \subseteq \mathcal{M}'$.

Proof: Suppose for the sake of the contradiction that it is not the case that for every maximal requirement consistent subset \mathcal{M}' of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}))$ there is some maximal consistent subset \mathcal{M} of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}_n))$ such that $\mathcal{M} \subseteq \mathcal{M}'$. That is, suppose there is some maximal requirement consistent subset \mathcal{M}' of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}))$ for which every maximal consistent subset \mathcal{M} of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}_n))$ is such that $\mathcal{M} \not\subseteq \mathcal{M}'$. Since \mathcal{M}' is requirement consistent, there must be some maximal consistent subset \mathcal{M}^* of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}_n))$ such that $\mathcal{M}' \cup \mathcal{M}^*$ is consistent. It follows from **Observation*** that \mathcal{M}^* is also a requirement consistent subset of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}))$. Thus, since \mathcal{M}' is also requirement consistent subset of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}))$ by stipulation, $\mathcal{M}^* \cup \mathcal{M}'$ is a requirement consistent subset of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}))$. Clearly $\mathcal{M}' \subseteq (\mathcal{M}^* \cup \mathcal{M}')$, and, since $\mathcal{M}^* \not\subseteq \mathcal{M}'$, $\mathcal{M}^* \cup \mathcal{M}' \neq \mathcal{M}'$. But this contradicts our assumption that \mathcal{M}' is a maximal requirement consistent subset of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}))$, because we have another requirement consistent subset $\mathcal{M}^* \cup \mathcal{M}'$, such that $\mathcal{M}' \subseteq (\mathcal{M}^* \cup \mathcal{M}') \subseteq \mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}))$. Therefore it must be that for every maximal requirement consistent subset \mathcal{M}' of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}))$ there is some maximal consistent subset \mathcal{M} of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}_n))$ such that $\mathcal{M} \subseteq \mathcal{M}'$.

Observation 3: Where $\langle \mathcal{R}, <, n \rangle$ is any threshold reasoning context, if $\langle \mathcal{R}, <, n \rangle \vdash \text{Must}(X)$ and $\langle \mathcal{R}, <, n \rangle \vdash \text{Ought}(Y)$, then it is not the case that $X \vdash \neg Y$.

Proof: If $\langle \mathcal{R}, <, n \rangle \vdash \text{Ought}(Y)$, then $\mathcal{M} \vdash Y$ for some maximal requirement consistent subset \mathcal{M} of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}))$. Since \mathcal{M} is requirement consistent, $\mathcal{M} \cup \mathcal{M}'$ is consistent for some maximal consistent subset \mathcal{M}' of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}_n))$. If $\langle \mathcal{R}, <, n \rangle \vdash \text{Must}(X)$, then $\mathcal{M}' \vdash X$. Now suppose for contradiction that $X \vdash \neg Y$. Then, by the transitivity of \vdash , $\mathcal{M}' \vdash \neg Y$. But then $\mathcal{M} \vdash Y$ and $\mathcal{M}' \vdash \neg Y$, contradicting our assumption that $\mathcal{M} \cup \mathcal{M}'$ is consistent. Therefore it is not the case that $X \vdash \neg Y$.

Observation 4: Where $\langle \mathcal{R}, <, n \rangle$ is any threshold reasoning context, if $\langle \mathcal{R}, <, n \rangle \vdash \text{Must}(X)$ and $\langle \mathcal{R}, <, n \rangle \vdash \text{Ought}(X)$.

Proof: If $\langle \mathcal{R}, <, n \rangle \vdash \text{Must}(X)$, then $\mathcal{M}' \vdash X$ for every maximal consistent subset \mathcal{M}' of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}_n))$. Let \mathcal{M}' be some such maximal consistent subset of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}_n))$. It follows from **Observation*** that \mathcal{M}' is also a

requirement consistent subset of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}))$. Either \mathcal{M}' is a maximal requirement consistent subset of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}))$ or not. If it is, then, since $\mathcal{M}' \vdash X$, $\langle \mathcal{R}, <, n \rangle \sim \text{Ought}(X)$. If not, then there is some maximal requirement consistent subset \mathcal{M}^* of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}))$ such that $\mathcal{M}' \subset \mathcal{M}^*$. In which case, since $\mathcal{M}' \subset \mathcal{M}^*$ and $\mathcal{M}' \vdash X$, it follows by the monotonicity of first order logic that $\mathcal{M}^* \vdash X$ and therefore $\langle \mathcal{R}, <, n \rangle \sim \text{Ought}(X)$. Thus $\langle \mathcal{R}, <, n \rangle \sim \text{Ought}(X)$.

Observation 5: Where $\langle \mathcal{R}, <, n \rangle$ is any threshold reasoning context, if $\langle \mathcal{R}, <, n \rangle \sim \text{Ought}(X)$, then $\langle \mathcal{R}, <, n \rangle \sim P(X)$.

Proof: If $\langle \mathcal{R}, <, n \rangle \sim \text{Ought}(X)$ then $\mathcal{M}' \vdash X$ for some maximal requirement consistent subset \mathcal{M} of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}))$. Since \mathcal{M} is requirement consistent, there is some maximal consistent subset \mathcal{M}' of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}_n))$ such that $\mathcal{M} \cup \mathcal{M}'$ is consistent. Thus there is some maximal consistent subset \mathcal{M}' of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}_n))$ such that $\mathcal{M}' \not\vdash \neg X$. Therefore it is not the case that $\langle \mathcal{R}, <, n \rangle \sim \text{Must}(X)$. Thus $\langle \mathcal{R}, <, n \rangle \sim P(X)$.

A3. Dual-Role Reasoning Contexts

Several proofs below invoke the following intermediate observations, which I again prove separately.

Observation[†]: Suppose that $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle$ is a dual-role reasoning context such that $\mathcal{D} \subseteq \mathcal{J}$. Every maximal consistent subset \mathcal{M} of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{D}))$ is also a requirement consistent subset \mathcal{M} of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{J}))$.

Proof: First, note that for any $r \in \text{Undefeated}_{<}(\mathcal{D})$, $r \in \mathcal{D}$ and therefore, since $\mathcal{D} \subseteq \mathcal{J}$, $r \in \mathcal{J}$. Second, for any $r \in \text{Undefeated}_{<}(\mathcal{D})$, by definition there is no other $r' \in \mathcal{R}$ such that (i) $r < r'$ and (ii) $\mathcal{C}(r') \vdash \neg \mathcal{C}(r)$. Thus for any $r \in \text{Undefeated}_{<}(\mathcal{D})$, $r \in \text{Undefeated}_{<}(\mathcal{J})$. So $\text{Undefeated}_{<}(\mathcal{D}) \subseteq \text{Undefeated}_{<}(\mathcal{J})$ and $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{D})) \subseteq \mathcal{C}(\text{Undefeated}_{<}(\mathcal{J}))$. Thus for any maximal consistent subset $\mathcal{M} \subseteq \mathcal{C}(\text{Undefeated}_{<}(\mathcal{D}))$, $\mathcal{M} \subseteq \mathcal{C}(\text{Undefeated}_{<}(\mathcal{J}))$. Clearly any maximal consistent subset $\mathcal{M} \subseteq \mathcal{C}(\text{Undefeated}_{<}(\mathcal{D}))$ is also requirement consistent. Therefore any maximal consistent subset \mathcal{M} of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{D}))$ is a requirement consistent of subset of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{J}))$.

Observation^{††}: Suppose that $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle$ is a dual-role reasoning context such that $\mathcal{D} \subseteq \mathcal{J}$. For every maximal requirement consistent subset \mathcal{M}'

of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{J}))$, there is some maximal consistent subset \mathcal{M} of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{D}))$ such that $\mathcal{M} \subseteq \mathcal{M}'$.

Proof: Assume for the sake of contradiction that it is not the case that for every maximal requirement consistent subset \mathcal{M}' of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{J}))$, there is some maximal consistent subset \mathcal{M} of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{D}))$ such that $\mathcal{M} \subseteq \mathcal{M}'$. That is, suppose there is some maximal requirement consistent subset \mathcal{M}' of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{J}))$ for which every maximal requirement consistent subset \mathcal{M} of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{D}))$ is such that $\mathcal{M} \not\subseteq \mathcal{M}'$. Since \mathcal{M}' is requirement consistent, $\mathcal{M}' \cup \mathcal{M}^*$ is consistent for some maximal consistent subset \mathcal{M}^* of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{D}))$. From **Observation**[†] it follows that \mathcal{M}^* is also a requirement consistent subset of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{J}))$. Thus, since \mathcal{M}' is also a requirement consistent subset of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{J}))$ by stipulation, $\mathcal{M}^* \cup \mathcal{M}'$ is a requirement consistent subset of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{J}))$. Clearly $\mathcal{M}' \subseteq (\mathcal{M}^* \cup \mathcal{M}')$, and, since $\mathcal{M}^* \not\subseteq \mathcal{M}'$, $\mathcal{M}^* \cup \mathcal{M}' \neq \mathcal{M}'$. But this contradicts our assumption that \mathcal{M}' is a maximal requirement consistent subset of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{J}))$, since there is another set $\mathcal{M}^* \cup \mathcal{M}'$ such that $\mathcal{M}' \subseteq (\mathcal{M}^* \cup \mathcal{M}') \subseteq \mathcal{C}(\text{Undefeated}_{<}(\mathcal{J}))$. Therefore it must be that for any maximal requirement consistent subset \mathcal{M}' of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{J}))$ there is some maximal consistent subset \mathcal{M} of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{D}))$ such that $\mathcal{M} \subseteq \mathcal{M}'$.

Observation 6: Where $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle$ is any dual – role reasoning context, if $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash \text{Must}(X)$ and $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash \text{Ought}(Y)$, then it is not the case that $X \vdash \neg Y$.

Proof: If $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash \text{Ought}(Y)$, then $\mathcal{M} \vdash Y$ for some maximal requirement consistent subset \mathcal{M} of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{J}))$. Since \mathcal{M} is requirement consistent, $\mathcal{M} \cup \mathcal{M}'$ is consistent for some maximal consistent subset \mathcal{M}' of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{D}))$. If $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash \text{Must}(X)$, then $\mathcal{M}' \vdash X$. Now suppose for contradiction that $X \vdash \neg Y$. By the transitivity of \vdash , $\mathcal{M}' \vdash \neg Y$. But then $\mathcal{M} \vdash Y$ and $\mathcal{M}' \vdash \neg Y$, contradicting our assumption that $\mathcal{M} \cup \mathcal{M}'$ is consistent. Therefore it is not the case that $X \vdash \neg Y$.

Observation 7: Suppose that $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle$ is a dual – role reasoning context such that $\mathcal{D} \subseteq \mathcal{J}$. If $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash \text{Must}(X)$, then $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash \text{Ought}(X)$.

Proof: If $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash \text{Must}(X)$, then $\mathcal{M} \vdash X$ for all maximal consistent subsets \mathcal{M} of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{D}))$. Let \mathcal{M}' be some such maximal consistent subset of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{D}))$. It follows from **Observation**[†] that \mathcal{M}' is a requirement consistent subset of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{J}))$. Either \mathcal{M}' is a maximal requirement consistent

subsets of or it is not. If it is, then $\mathcal{M}' \vdash X$ and thus $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash \text{Ought}(X)$. If it is not, then there is some maximal requirement consistent subset \mathcal{M}^* of $\mathcal{C}(\text{Undeclared}_z(\mathcal{J}))$ such that $\mathcal{M}' \subset \mathcal{M}^*$. In which case, since $\mathcal{M}' \vdash X$ and $\mathcal{M}' \subset \mathcal{M}^*$, it follows from the monotonicity of first order logic that $\mathcal{M}^* \vdash X$ and thus $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash \text{Ought}(X)$. Thus $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash \text{Ought}(X)$.

Observation 8: Where $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle$ is any dual-role reasoning context. If $\langle \mathcal{R}, \mathcal{J}, \mathcal{D}, < \rangle \vdash \text{Ought}(X)$, then $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash P(X)$

Proof: If $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash \text{Ought}(X)$, then $\mathcal{M} \vdash X$ for some maximal requirement consistent subject of \mathcal{M} of $\mathcal{C}(\text{Undeclared}_z(\mathcal{J}))$. Since \mathcal{M} is requirement consistent, there is some maximal consistent subsets \mathcal{M}' of $\mathcal{C}(\text{Undeclared}) < (\mathcal{D})$ such that $\mathcal{M} \cup \mathcal{M}'$ is consistent. Thus there is some maximal consistent subset \mathcal{M}' of $\mathcal{C}(\text{Undeclared}_z(\mathcal{D}))$ such that $\mathcal{M}' \not\vdash \neg X$. Therefore it is not the case that $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash \text{Must}(\neg X)$. Thus $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash P(X)$.

A4. Ought-Requirement Agglomeration

For any reasoning context Δ , the principle of ought-requirement agglomeration applies just in case if $\Delta \vdash \text{Must}(X)$ and $\Delta \vdash \text{Ought}(Y)$, then $\Delta \vdash \text{Ought}(X \wedge Y)$. This principle holds some appeal in its own right. Suppose that in a given reasoning context Δ I am required to go to the beach, so that $\Delta \vdash \text{Must}(B)$. Further suppose that in that same reasoning context I ought to buy myself an ice-cream, so that $\Delta \vdash \text{Ought}(I)$. Ought-requirement agglomeration supports the further conclusion $\Delta \vdash \text{Ought}(B \wedge I)$, that I ought to go to the beach and buy myself an ice-cream, which seems plausible. However, ought-requirement agglomeration is particularly attractive in light of the fact that, when combined with the closure of oughts under logical consequence, it preserves other attractive patterns of inference (cf. Horty 2012: 87–89).

Two scenarios illustrate the attractiveness of combining the principle of ought-requirement agglomeration with the closure of oughts under logical consequence. First suppose that in a given reasoning context Δ , John ought to fight in the army or serve his country in some other way (by serving as a hospital volunteer, for instance), so that we have that $\Delta \vdash \text{Ought}(F \vee S)$. John is not required to flight or serve, but he ought to in the circumstances. Now suppose that John is required not to flight because his religion forbids it, so that $\Delta \vdash \text{Must}(\neg F)$. From ought-requirement agglomeration, we can conclude that $\Delta \vdash \text{Ought}((F \vee S) \wedge \neg F)$. Since $(F \vee S) \wedge \neg F \vdash S$, the closure of ought under

logical consequence then allows us to reach the conclusion $\Delta \vdash \text{Ought}(S)$, that John ought to serve his country in some other way, which seems to be the correct result in this reasoning context.

Second, suppose that in a different given reasoning context Δ' a student has two sets of homework he could finish this evening: his math homework, which is due first thing in the morning, and his English homework, which is not due until the following day. Because it is due first thing, he student is required to finish his math, so that $\Delta' \vdash \text{Must}(M)$. Suppose, however, that he really ought to finish his English homework to free up time tomorrow evening, so that $\Delta' \vdash \text{Ought}(E)$. Suppose further that $M \wedge E \vdash \neg H$. If the student finishes both his math and English homework, then he will have no more homework to complete at the end of evening. Ought-requirement agglomeration allows us to conclude that $\Delta' \vdash \text{Ought}(M \wedge E)$. The student ought to finish his math and his English homework. Ought Closure then allows us to conclude $\Delta' \vdash \text{Ought}(\neg H)$, that the student ought to have no more homework to complete at the end of the evening. Again, this seems like the correct result in this reasoning context.

To begin with, we note that both threshold and dual-role reasoning contexts preserve the closure of ought under logical consequence.

Observation 9: Where $\langle \mathcal{R}, <, n \rangle$ is any threshold reasoning context, if $\langle \mathcal{R}, <, n \rangle \vdash \text{Ought}(X)$ and $X \vdash Y$ then $\langle \mathcal{R}, <, n \rangle \vdash \text{Ought}(Y)$.

Proof: If $\langle \mathcal{R}, <, n \rangle \vdash \text{Ought}(X)$, then $\mathcal{M} \vdash X$ for some maximal requirement consistent subset \mathcal{M} of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}))$. So $\mathcal{M} \vdash X$ and $X \vdash Y$ and thus, by the transitivity of \vdash , $\mathcal{M} \vdash Y$. Therefore $\langle \mathcal{R}, <, n \rangle \vdash \text{Ought}(Y)$.

Observation 10: Where $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle$ is any dual-role reasoning context, if $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash \text{Ought}(X)$ and $X \vdash Y$ then $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash \text{Ought}(Y)$.

Proof: If $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash \text{Ought}(X)$, then $\mathcal{M} \vdash X$ for some maximal requirement consistent subset \mathcal{M} of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{J}))$. So $\mathcal{M} \vdash X$ and $X \vdash Y$ and thus, by the transitivity of \vdash , $\mathcal{M} \vdash Y$. Therefore $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash \text{Ought}(Y)$.

Both threshold and dual-role reasoning contexts also both preserve ought-requirement agglomeration.

Observation 11: Where $\langle \mathcal{R}, <, n \rangle$ is any threshold reasoning context, if $\langle \mathcal{R}, <, n \rangle \vdash \text{Must}(X)$ and $\langle \mathcal{R}, <, n \rangle \vdash \text{Ought}(Y)$ then $\langle \mathcal{R}, <, n \rangle \vdash \text{Ought}(X \wedge Y)$.

Proof: If $\langle \mathcal{R}, <, n \rangle \vdash \text{Ought}(Y)$, then $\mathcal{M}' \vdash Y$ for some maximal requirement consistent subset \mathcal{M}' of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}))$. It follows from **Observation**** that there is some maximal consistent subset \mathcal{M} of $\mathcal{C}(\text{Undefeated}_{<}(\mathcal{R}_n))$ such that $\mathcal{M} \subseteq \mathcal{M}'$. If $\langle \mathcal{R}, <, n \rangle \vdash \text{Must}(X)$, then $\mathcal{M} \vdash X$. Since $\mathcal{M} \subseteq \mathcal{M}'$ and $\mathcal{M} \vdash X$ it follows by the monotonicity of first order logic that $\mathcal{M}' \vdash X$. So $\mathcal{M}' \vdash X$ and $\mathcal{M}' \vdash Y$ and therefore $\mathcal{M}' \vdash X \wedge Y$. Therefore $\langle \mathcal{R}, <, n \rangle \vdash \text{Ought}(X \wedge Y)$.

Observation 12: Suppose that $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle$ is a dual-role reasoning context such that $D \subseteq \mathcal{J}$. If $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash \sim \text{Must}(X)$ and $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash \text{Ought}(Y)$, then $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash \text{Ought}(X \wedge Y)$.

Proof: If $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash \text{Ought}(Y)$, then $\mathcal{M}' \vdash Y$ for some maximal requirement consistent subset \mathcal{M}' of $\mathcal{C}(\text{Undeclared}_<(\mathcal{J}))$. It follows from **Observation^{††}** that there is some maximal consistent subset \mathcal{M} of $\mathcal{C}(\text{Undeclared}_<(\mathcal{D}))$ such that $\mathcal{M} \subseteq \mathcal{M}'$. If $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash \sim \text{Must}(X)$, then $\mathcal{M} \vdash X$. Since $\mathcal{M} \subseteq \mathcal{M}'$ and $\mathcal{M} \vdash X$, it follows by the monotonicity of first order logic that $\mathcal{M}' \vdash X$. So $\mathcal{M}' \vdash X$ and $\mathcal{M}' \vdash Y$ and therefore $\mathcal{M}' \vdash X \wedge Y$. Thus $\langle \mathcal{R}, \mathcal{D}, \mathcal{J}, < \rangle \vdash \text{Ought}(X \wedge Y)$.

Acknowledgements

I would like to thank my referees for very helpful comments on earlier versions of the paper. I am particularly grateful to Shyam Nair, who not only provided me with extensive, thoughtful comments on previous versions of the manuscript, but corrected several errors in previous proofs and encouraged me to make important amendments. All remaining errors are my own.

References

- Bedke, Matt (2011). Passing the Deontic Buck. *Oxford Studies in Metaethics*, 6, 128.
- Broome, John (2016). A Linguistic Turn in the Philosophy of Normativity? *Analytic Philosophy*, 57(1), 1–14.
- Chisholm, Roderick M. (1974). Practical Reason and the Logic of Requirement. In Stephan Korner (Ed.), *Practical Reason* (1–17). Oxford University Press.
- Dancy, Jonathan (1993). *Moral Reasons*. Blackwell.
- Dancy, Jonathan (2004). Enticing Reasons. In R. Jay Wallace, Philip Pettit, Samuel Scheffler, and Michael Smith (Eds.), *Reason and Value: Themes from the Moral Philosophy of Joseph Raz* (91–118). Clarendon Press.
- van Fraassen, Bas C. (1973). Values and the Heart's Command. *Journal of Philosophy*, 70(1), 5–19.
- Gert, Joshua (2004). *Brute Rationality: Normativity and Human Action*. Cambridge University Press.
- Gert, Joshua (2007). Normative Strength and the Balance of Reasons. *Philosophical Review*, 116(4), 533–562.
- Gert, Joshua (2012). Moral Worth, Supererogation, and the Justifying/Requiring Distinction. *Philosophical Review*, 121(4), 611–618.
- Gert, Joshua (2014). Perform a Justified Option. *Utilitas*, 26(2), 206–217.
- Goble, Lou (2013). Prima Facie Norms, Normative Conflicts, and Dilemmas. In Dov Gabbay, Xavier Parent, Ron van der Mayden, and Leendert van der Torre (Eds.), *Handbook of Deontic Logic and Normative Systems* (241–352). College Publications.

- Greenspan, Patricia (2005). Asymmetrical Practical Reasons. In M. E. Reicher and J. C. Marek (Eds.) *Experience and Analysis: Proceedings of the 27th International Wittgenstein Symposium* (387–394). Öbv & hpt.
- Greenspan, Patricia (2007). Practical Reasons and Moral 'Ought'. In Russ Shafer-Landau (Ed.), *Oxford Studies in Metaethics* (Vol. 2, 172–194). Oxford University Press.
- Greenspan, Patricia (2010). Making Room for Options: Moral Reasons, Imperfect Duties, and Choice. *Social Philosophy and Policy*, 27(2), 181–205.
- Hampton, Jean (1998). *The Authority of Reason*. Cambridge University Press.
- Hansen, Jörg (2004). Problems and Results for Logics about Imperatives. *Journal of Applied Logic*, 2(1), 39–61.
- Hansen, Jörg (2005). Conflicting Imperatives and Dyadic Deontic Logic. *Journal of Applied Logic*, 3(3–4), 484–511.
- Heuer, Ulrike (2004). Raz on Values and Reasons. In R. Jay Wallace, Philipp Pettit, Samuel Scheffler, and Michael Smith (Eds.), *Reason and Value: Themes from the Philosophy of Joseph Raz* (129–152). Oxford University Press.
- Heyd, David (1982). *Supererogation: Its Status in Ethical Theory*. Cambridge University Press.
- Horgan, Terry, and Mark Timmons (2010). Untying a Knot from the Inside Out: Reflections on the 'Paradox' of Supererogation. *Social Philosophy and Policy*, 27(2), 29–63.
- Horty, John (2003). Reasoning with Moral Conflicts. *Noûs*, 37(4), 557–605.
- Horty, John (2012). *Reasons as Defaults*. Oxford University Press.
- Kagan, Shelly (1989). *The Limits of Morality*. Oxford University Press.
- Little, Margaret (2013). In Defence of Non-Deontic Reasons. In David Bakhurst, Margaret Little, and Brad Hooker (Eds.), *Thinking About Reasons: Themes From the Philosophy of Jonathan Dancy* (112–136). Oxford University Press.
- Macnamara, Coleen and Margaret Little (2017). For Better or Worse: Commendatory Reasons and Latitude. In Mark Timmons (Ed.) *Oxford Studies in Normative Ethics* (Vol. 7, 138–160). Oxford University Press.
- Marcus, Ruth Barcan (1980). Moral Dilemmas and Consistency. *Journal of Philosophy*, 77(3), 121–136.
- McNamara, Paul (1996). Must I Do What I Ought (Or Will the Least I Can Do Do)? In Mark Brown and José Carmo (Eds.), *Deontic Logic, Agency and Normative Systems*, (154–173). Springer-Verlag.
- Nair, Shyam (2016). Conflicting Reasons, Unconflicting 'Oughts.' *Philosophical Studies*, 173(3), 629–663.
- Nair, Shyam and John Horty (2018). The Logic of Reasons. In Daniel Star (Ed.) *The Oxford Handbook of Reasons and Normativity* (67–83). Oxford University Press.
- Parfit, Derek (2011). *On What Matters*. Oxford University Press.
- Portmore, Douglas W. (2013). Perform Your Best Option. *Journal of Philosophy*, 110(8), 436–459.
- Portner, Paul (2009). *Modality*. Oxford University Press.
- Raz, Joseph (1975). Permissions and Supererogation. *American Philosophical Quarterly*, 12(2), 161–168.
- Raz, Joseph (1999a). *Practical Reason and Norms* (2nd ed.). Oxford University Press.
- Raz, Joseph (1999b). *Engaging Reason: On the Theory of Value and Action*. Oxford University Press.
- Scanlon, T. M. (1998). *What We Owe to Each Other*. Harvard University Press.
- Scanlon, T. M. (2014). *Being Realistic About Reasons*. Oxford University Press.

- Schroeder, Mark (2007). *Slaves of the Passions*. Oxford University Press.
- Schroeder, Mark (2011). Holism, Weight and Undercutting. *Noûs*, 45(2), 328–344.
- Silk, Alex (2015). What Normative Terms Mean and Why It Matters for Ethical Theory. In Mark Timmons (Ed.) *Oxford Studies in Normative Ethics* (Vol. 5, 296–325). Oxford University Press.
- Skorupski, John (2010). *The Domain of Reasons*. Oxford University Press.
- Snedegar, Justin (2016). Reasons, Oughts, and Requirements. *Oxford Studies in Metaethics* (Vol. 11, 155–181). Oxford University Press.
- Tucker, Dustin (2018). Variable Priorities and Exclusionary Reasons in Input/Output Logic. *Journal of Philosophical Logic*, 47(6), 947–964.
- Williams, Bernard (1981). *Moral Luck: Philosophical Papers 1973–1980*. Cambridge University Press.